



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

1/1

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/052,435	01/23/2002	Michael Kagan	3891-0102P	1934
7590	01/17/2006		EXAMINER	
Dr. Mark Friedman Ltd. c/o Discovery Dispatch 9003 Florin Way Upper Marlboro, MD 20772			MADAMBA, GLENFORD J	
			ART UNIT	PAPER NUMBER
			2151	

DATE MAILED: 01/17/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)
	10/052,435	KAGAN ET AL.
	Examiner Glenford Madamba	Art Unit 2151

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE ____ MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 19 October 2005.
- 2a) This action is FINAL. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1-40 is/are pending in the application.
- 4a) Of the above claim(s) ____ is/are withdrawn from consideration.
- 5) Claim(s) ____ is/are allowed.
- 6) Claim(s) 1-40 is/are rejected.
- 7) Claim(s) ____ is/are objected to.
- 8) Claim(s) ____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on 19 October 2005 is/are: a) accepted or b) objected to by the Examiner. Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a). Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. ____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|---|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. ____. |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date ____. | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| | 6) <input type="checkbox"/> Other: ____. |

DETAILED ACTION

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

2. Claims 1-8, 10-14, 17, 22, 24-26, 28-31, 34, 37 are rejected under 35 U.S.C. 102(e) as being anticipated by Tzeng et al, U.S. Patent 6,912, 604.

3. As per claims 1 and 28, Tzeng discloses an interface adapter for a packet network, comprising:

a memory interface, for coupling to a memory (HCA_12 coupled to Memory Controller_24 and System Memory_26) [Fig. 1] (System/External Memory_48) [Fig. 2];
a first plurality of execution engines (Submodules 68a, 68b, 68c, and 6d of Transport Service Module_42) [Fig. 1] [col 2, lines 37-45], coupled to the host interface so as to read from the memory work items (WQEs) corresponding to messages to be sent over the network, and to generate gather entries (descriptors) defining packets to

be transmitted over the network responsive to the work items [col 6, line 41 – col 7, line 4] [col 8, lines 40-44] [col 9, lines 29-32]; and

 a scheduling processor, coupled to assign the work items to the execution engines for generation of the gather entries (Pre-Link Module_40) [Fig. 2] [col 2, lines 37-45] [col 4, lines 1-14] [col 4, line 48 – col 5, line 12] [col 8, lines 35-39 & 48-61[[col 9, lines 25-28];

 a second plurality of gather engines, which are adapted to generate the packets responsive to the gather entries (Post-Link Module_44) [Fig. 2] [col 2, lines 37-45] [col 7, lines 5-18] [col 8, lines 45-48] [col 9, lines 32-35]; and

 switching circuitry (Link-layer arbitration tables and databases) [Fig. 2], coupling the execution engines to the gather engines so as to submit the gather entries to the gather engines for generation of the packets responsive thereto (HCA comprising Pre-link, Transport Service, and Post-link module communicatively coupled to each other) [col 9, lines 24-35].

4. As per claim 2, Tzeng discloses an adapter according to claim 1, wherein each of the work items belongs to a respective transport service instance among multiple transport service instances served by the adapter, and wherein the scheduling processor is adapted to assign the work items to the execution engines by selecting the respective transport service instances for service [col 6, lines 16-40].

5. As per claims 3 and 29, Tzeng discloses an adapter according to claim 2, wherein the transport service instances comprise queue pairs, and wherein the work items comprise work queue elements [col 5, lines 34-59] [col 6, lines 41-55].

6. As per claims 4 and 30, Tzeng, discloses an adapter according to claim 2, wherein for each of the transport service instances, the respective work items are maintained in a list in the memory (Create L4 info in ext-memory) [Fig. 2], and wherein the execution engines are configured so as to generate in succession the gather entries corresponding to the work items in the list maintained for the transport service instance selected by the scheduling processor [col 8, lines 40-44] [col 9, lines 29-35] [Fig. 2].

7. As per claim 5, Tzeng discloses an adapter according to claim 2, wherein the scheduling processor is adapted to maintain a third plurality of scheduling queues (Queue Pair FIFO_62) [Fig. 2] in which are entered the transport service instances to which the work items belong, the scheduling queues having respective heads and tails, and to select the instances from the heads of the queues for assignment to the execution engines [col 5, lines 34-42] [col 6, line 41 – col 7, line 4].

8. As per claim 6, Tzeng discloses an adapter according to claim 5, wherein the scheduling queues are associated with respective classes of service provided by the adapter, and wherein the transport service instances are assigned to the scheduling

queues according to the classes of services to which the corresponding transport services belong [col 5, line 61 – col 6, line 40].

9. As per claims 7 and 31, Tzeng discloses an adapter according to claim 6, wherein the scheduling processor is adapted to determine the scheduling queues from which the transport service instances are to be assigned to the execution engines in accordance with a scheduling policy relating to the respective classes of service of the scheduling queues [col 5, line 61 – col 6, line 40].

10. As per claim 8, Tzeng discloses an adapter according to claim 7, wherein the switching circuitry is coupled to arbitrate among the execution engines so as to submit the gather entries to the gather engines in an order responsive to the classes of service [col 5, lines 13-33].

11. As per claims 10 and 34, Tzeng discloses an adapter according to claim 5, wherein context information regarding each of the transport service instances is recorded in the memory, and wherein the scheduling processor is adapted to maintain the scheduling queues by directing pointers in the context information of each of the instances, except the instances at the tails of the scheduling queues, to point to the context information of succeeding instances in the queues [Tzeng: col 6, lines 41-65].

12. As per claim 11, Tzeng discloses an adapter according to claim 2, wherein the transport service instances are assigned to respective classes of service provided by the adapter, and wherein the scheduling processor is adapted to select the transport service instances for service by the execution engines responsive to the assigned classes of service [col 4, line 63 – col 5, line33] [col 6, lines 16-47].

13. As per claim 12, Tzeng discloses an adapter according to claim 11, wherein the switching circuitry is coupled to arbitrate among the execution engines so as to submit the gather entries to the gather engines in an order responsive to the classes of service [col 5, lines 13-33].

14. As per claim 13, Tzeng discloses an adapter according to claim 1, wherein the work items belong to different transport service instances, which are associated with respective classes of service, and wherein the scheduling processor is adapted to select the work items to be assigned to the execution engines responsive to the classes of service [col 4, line 63 – col 5, line33] [col 6, lines 16-47].

15. As per claim 14, Tzeng discloses an adapter according to claim 13, wherein the switching circuitry is coupled to arbitrate among the execution engines so as to submit the gather entries to the gather engines in an order responsive to the classes of service [col 5, lines 13-33].

16. As per claim 17, Tzeng discloses an adapter according to claim 1, wherein the work items comprise descriptors indicating data to be read from the memory for inclusion in the packets, and wherein the gather engines are coupled to read the data from the memory by direct memory access (DMA) [Fig. 2].

17. As per claim 22, Tzeng discloses an adapter according to claim 1, wherein the work items belong to different transport service instances, and wherein the scheduling processor comprises a microprocessor (embedded processor_80), which is programmed by software code to select the transport service instances for assignment to the execution engines, and wherein the switching circuitry comprises [col 4, line 63 – col 5, line 13]:

one or more registers, containing respective weight factors associated with the execution engines [col 5, lines 13-33] [col 8, lines 11-24] [col 9, lines 24-35]; and

one or more hardware logic circuits, coupled to arbitrate among the execution engines to submit the gather entries to each of the gather engines responsive to the weight factors [col 5, lines 13-33] [col 8, lines 11-24] [col 9, lines 24-35].

18. As per claim 24, Tzeng discloses an adapter according to claim 22, wherein the one or more hardware logic circuits comprise a respective arbiter circuit connected to each one of the second plurality of gather engines, and wherein the arbiter circuit is coupled to arbitrate among the execution engines whose gather entries are to be submitted to the one of the gather engines connected thereto [col 5, lines 13-33] [col 8,

lines 11-24].

19. As per claim 25, Tzeng discloses an adapter according to claim 24, wherein the scheduling processor is adapted to determine which of the execution engines are to submit the gather entries to each of the gather engines, and to assign the execution engines to submit the gather entries accordingly, and wherein the arbiter circuit connected to each of the gather engines is coupled to arbitrate among the execution engines assigned thereto [col 6, line 41 – col 7, line 45] [col 8, lines 11-24].

20. As per claim 26, Tzeng discloses an adapter according to claim 1, wherein the first plurality is greater than the second plurality [Fig. 2].

21. As per claim 37, Tzeng discloses an adapter according to claim 28, wherein at least some of the messages comprise data to be read from the memory and sent to a recipient via the network, and wherein the work items indicate the data to be sent [col 4, lines 28-39][Fig.1].

Claim Rejections - 35 USC § 103

1. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

2. Claims 9, 27, and 32 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tzeng in view of Pettey et al (hereinafter Pettey), U.S. Patent 6,594,712.

3. As per claims 9 and 32, Tzeng in view of Pettey discloses an adapter according to claim 5, wherein the scheduling processor is further adapted to enter the transport service instances at the tails of the scheduling queues to which they are assigned when work items belonging to the transport service instances are written to the memory.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order. One aspect of the present invention provides a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the

determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 5, he does not explicitly disclose wherein the scheduling processor is further adapted to enter the transport service instances *at the tails of the scheduling queues* to which they are assigned when work items belonging to the transport service instances are written to the memory. This feature is disclosed by Pettey in his invention of an Infiniband channel adapter for performing direct data transfers between a PCI bus and Infiniband link without double-buffering the data in system memory [Fig. 23] [col 23, lines 21-47].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of entering the transport service instances at the tails of the scheduling queues to which they are assigned when work items are written to the memory, as disclosed by Pettey for the motivation of reducing host memory bandwidth consumption [col 25, lines 17-26].

4. As per claim 27, Tzeng in view of Petty discloses an adapter according to claim 1, wherein the packets generated by the gather engines are transmitted over the network at a given transmission speed of the network, and wherein the first plurality and the

second plurality are chosen so as to be capable of generating the packets at a speed greater than the given transmission speed.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 1 comprising a memory interface, a first plurality of execution engines to generate gather entries (descriptors) defining packets to be transmitted over the network responsive to the work items, a scheduling processor, and switching circuitry to submit the gather entries to the gather engines for generation of the packets responsive thereto; he does not explicitly disclose the adapter wherein packets generated by the

gather engines are transmitted over the network at a given transmission speed of the network, and wherein the first plurality and the second plurality are chosen so as to be capable of generating the packets at a speed greater than the given transmission speed.

This feature is disclosed by Pettey in his invention of an Infiniband channel adapter for performing direct data transfers between a PCI bus and Infiniband link without double-buffering the data in system memory [Abstract]. Pettey teaches that the minimum data transfer speed specified by the IBA is 2.5 Gbps, but can also go as high as 10-30 Gbps between IB-capable computers and I/O units [col 1, lines 19-30] [col 2, lines 22-43].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of transmitting packets generated by the gather engines over the network at a given transmission speed of the network (default/minimum required speed of 2.5 Gbps), and choosing the first plurality and the second plurality in order to generate the packets at a speed greater than the given transmission speed, as disclosed by Pettey for the motivation of reducing host memory bandwidth consumption [col 25, lines 17-26] and maximizing the transmission speed [col 1, lines 28-39 & 54-65].

5. Claims 18 and 19 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tzeng in view of Applicant's Admitted Prior Art (AAPA).

6. As per claim 18, Tzeng in view AAPA discloses an adapter according to claim 17, wherein the messages comprise remote direct memory access (RDMA) requests, and wherein the work items are written to the memory by a host processor submitting the requests.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3]. .

While Tzeng discloses substantial features of the invention such as the adapter of claim 17 wherein the work items comprise descriptors indicating data to be read from the memory for inclusion in the packets, and wherein the gather engines are coupled to read the data from the memory by direct memory access, he does not explicitly disclose the adapter *wherein the messages comprise remote direct memory access (RDMA) requests*, and wherein the work items are written to the memory by a host processor submitting the requests.

This feature is disclosed by Applicant in the background disclosure for his invention. Applicant discloses that the messages comprise RDMA request/response messages for writing and/or reading data to/from memory [0005]. It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of messages comprising remote direct memory access (RDMA) requests, and wherein the work items are written to the memory by a host processor submitting the requests, as disclosed by AAPA for the motivation of data transfer efficiency and performing the well-known operation of injecting messages into a fabric in an IB network architecture [0005].

7. As per claim 19, Tzeng in view of AAPA discloses an adapter according to claim 17, wherein the messages comprise remote direct memory access (RDMA) responses, and

wherein the work items are written to the memory responsive to RDMA request packets received by the adapter via the network.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 17 wherein the work items comprise descriptors indicating data to be read from the memory for inclusion in the packets, and wherein the gather engines are coupled to read the data from the memory by direct memory access, he does not explicitly disclose the adapter wherein *the messages comprise remote direct memory access (RDMA)*

responses, and wherein the work items are written to the memory responsive to RDMA request packets received by the adapter via the network.

This feature is disclosed by Applicant in the background disclosure for his invention. Applicant discloses that the messages comprise RDMA request/response messages for writing and/or reading data to/from memory [0005]. It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of *messages comprising remote direct memory access (RDMA) responses, and wherein the work items are written to the memory responsive to RDMA request packets received by the adapter via the network*, as disclosed by AAPA, for the motivation of data transfer efficiency and performing the well-known operation of injecting messages into a fabric in an IB network architecture [0005].

8. Claims 15, 20, 21, 23, 35, 36, and 38 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tzeng in view of Parthasarathy et al (hereinafter Parthasarathy), U.S. Patent 6,831,916.

9. As per claims 15 and 38, Tzeng in view of Parthasarathy discloses an adapter according to claim 13, and comprising an execution access *arbiter*, coupled between the execution engines and the memory interface so as to control an order of access to

the memory by the execution engines in reading the work items, responsive to the classes of service.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 13, wherein the work items belong to different transport service instances, which are associated with respective classes of service, and wherein the scheduling processor is adapted to select the work items to be assigned to the execution engines responsive to the classes of service, he does not explicitly disclose the adapter comprising an execution access *arbiter*, coupled between the execution engines and

the memory interface so as to control an order of access to the memory by the execution engines in reading the work items, responsive to the classes of service.

This limitation is disclosed by Parthasarathy in the disclosure for his invention of a host system that comprises one or more host-fabric adapters installed therein for connecting to a switched fabric of a data network [Abstract]. Referring to Figure 7, Parthasarathy discloses that the host interface 712 provides an interface to either an I/O bus 205 of a host system 130 (FIG. 3), or an I/O and memory controller 204 of a host system 130 (FIG. 4) for host transactions, *including controlling arbitration* and data/control multiplexing between different requesters, read and write transactions to the host system 130 and facilitating read completions [col 10, lines 6-12].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of the adapter comprising an execution access *arbiter*, coupled between the execution engines and the memory interface so as to control an order of access to the memory by the execution engines in reading the work items, responsive to the classes of service, as disclosed by Parthasarathy, for the motivation of increased performance and efficiency and optimizing NGIO/Infiniband functionality with minimal hardware investment [Parthasarathy: col 9, lines 37-43].

10. As per claims 20 and 35, Tzeng in view of Parthasarathy discloses an adapter according to claim 17, and comprising a data access arbiter, coupled between the gather engines and the memory interface so as to control an order of access to the memory by the gather engines for reading the data.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 13, wherein the work items belong to different transport service instances, which are associated with respective classes of service, and wherein the scheduling processor is adapted to select the work items to be assigned to the execution engines

responsive to the classes of service, he does not explicitly disclose the adapter comprising a data access arbiter, coupled between the gather engines and the memory interface so as to control an order of access to the memory by the gather engines for reading the data.

This limitation is disclosed by Parthasarathy in the disclosure for his invention of a host system that comprises one or more host-fabric adapters installed therein for connecting to a switched fabric of a data network [Abstract]. Referring to Figure 7, Parthasarathy discloses that the host interface 712 provides an interface to either an I/O bus 205 of a host system 130 (FIG. 3), or an I/O and memory controller 204 of a host system 130 (FIG. 4) for host transactions, *including controlling arbitration* and data/control multiplexing between different requesters, read and write transactions to the host system 130 and facilitating read completions [col 10, lines 6-12].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of the adapter comprising a data access arbiter, coupled between the gather engines and the memory interface so as to control an order of access to the memory by the gather engines for reading the data, as disclosed by Parthasarathy, for the motivation of increased performance and efficiency and optimizing NGIO/Infiniband functionality with minimal hardware investment [Parthasarathy: col 9, lines 37-43].

11. As per claims 21 and 36, Tzeng discloses an adapter according to claim 20, wherein the work items belong to different transport service instances, which are associated with respective classes of service, and wherein the data access arbiter is programmable so as to give priority in the order of access to one or more of the gather engines responsive to the classes of service [col 5, lines 13-33] [col 8, lines 11-24] [col 9, lines 24-35].

12. As per claim 23, Tzeng in view of Parthasarathy discloses an adapter according to claim 22, wherein the memory interface, execution engines, scheduling processor, gather engines and switching circuitry are comprised in a single integrated circuit chip.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets

having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 22, wherein the work items belong to different transport service instances, and wherein the scheduling processor comprises a microprocessor, which is programmed by software code to select the transport service instances for assignment to the execution engines, and wherein the switching circuitry comprises one or more registers, containing respective weight factors associated with the execution engines; and one or more hardware logic circuits, coupled to arbitrate among the execution engines to submit the gather entries to each of the gather engines responsive to the weight factors; he does not explicitly disclose the adapter wherein the memory interface, execution engines, scheduling processor, gather engines and switching circuitry are comprised in a single integrated circuit chip.

This limitation is disclosed by Parthasarathy in the disclosure for his invention of a host system that comprises one or more host-fabric adapters installed therein for connecting to a switched fabric of a data network [Abstract]. Parthasarathy expressly discloses that the invention is applicable for use with all types of data networks, I/O hardware adapters and *chipsets*, including follow-on *chip designs* which link together end stations and communication devices for data communications [col 3, lines 20-24].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of an adapter having components such as the memory interface, execution engines, scheduling processor, gather engines and switching circuitry comprised in a single integrated circuit chip, as disclosed by Parthasarathy, for the motivation of increased performance and efficiency and optimizing NGIO/Infiniband functionality with minimal hardware investment [Parthasarathy: col 9, lines 37-43].

13. Claim 33 is rejected under 35 U.S.C. 103(a) as being unpatentable over Tzeng in view of Pettey et al (hereinafter Pettey), and in further view of Parthasarathy et al (hereinafter Parthasarathy), U.S. Patent 6,831,916.

14. As per claim 33, Tzeng in view of Pettey and in further view of Parthasarathy discloses an adapter according to claim 32, wherein the scheduling processor is adapted to enter the transport service instance in the scheduling queues responsive to a host processor having written to a doorbell address of the adapter.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work

queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

Pettey discloses as his invention an Infiniband channel adapter for performing direct data transfers between a PCI bus and Infiniband link without double-buffering the data in system memory [Fig. 23] [col 23, lines 21-47]. As discussed for claim 32, Tzeng in view of Pettey discloses the adapter wherein the scheduling processor is adapted to enter the transport service instance in the scheduling queues. But while Tzeng and Pettey disclose substantial features of the invention, such as the adapter described by claims 28 and 32, neither Tzeng nor Pettey expressly discloses the adapter wherein the scheduling processor is adapted to enter the transport service instance (WQEs) in the scheduling queues *responsive to a host processor having written to a doorbell address of the adapter.*

This feature is disclosed by Parthasarathy in the disclosure of his invention of a host system that comprises one or more host-fabric adapters installed therein for connecting to a switched fabric of a data network [Abstract]. Referring to Figure 7,

Parthasarathy discloses that the host fabric adapter (120) includes a micro-controller subsystem (700) which itself contains one or more programmable DMA engines (Micro-Engine {ME}) utilized to build, send, receive, and acknowledge NGIO/Infiniband cells between the host memory (206) and a serial link, and special purpose hardware logic blocks such as host interface (712), address translation interface (714), a VI context memory interface (716), local bus interface (716) and a completion queue/*doorbell manager* interface (720)[col 9, lines 44-64] [col 10, 37-40] [Fig. 7].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify the combined invention of Tzeng in view of Pettey with the additional feature of the adapter wherein the scheduling processor is adapted to enter the transport service instance (WQEs) in the scheduling queues *responsive to a host processor having written to a doorbell address of the adapter*, as disclosed by Parthasarathy, for the motivation of increased performance and efficiency and optimizing NGIO/Infiniband functionality with minimal hardware investment [Parthasarathy: col 9, lines 37-43].

15. Claim 16 is rejected under 35 U.S.C. 103(a) as being unpatentable over Tzeng in view of Grun, U.S. Patent 6,272,591.

16. As per claim 16, Tzeng in view of Grun discloses an adapter according to claim 1, wherein the execution engines are adapted to generate multiple gather entries corresponding to a single one of the work items, *each of the gather entries defining no more than a single packet to be generated by one of the gather engines.*

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 1 comprising a memory interface, a first plurality of execution engines to generate gather entries (descriptors) defining packets to be transmitted over the network responsive to the work items, a scheduling processor, and switching circuitry to

submit the gather entries to the gather engines for generation of the packets responsive thereto; he does not explicitly disclose the adapter wherein the execution engines are adapted to generate multiple gather entries corresponding to a single one of the work items, *each of the gather entries defining no more than a single packet* to be generated by one of the gather engines.

This feature is disclosed by Grun in his invention of a RAID device for striping a data block across N disk drives. Grun teaches that in most RAID devices, a host computer sends an entire data block in one piece to the RAID controller. The RAID device receives a storage request from a host computer for the data block, and creates N virtual interface ("VI") queue pairs. The queue pairs form N virtual channels to the host computer. Further, the RAID device posts a descriptor to each of the queue pairs, with each descriptor referring to 1/Nth of the data block. Further, the RAID device receives 1/Nth of the data block over each of the virtual channels and writes each received 1/Nth data block to a different one of the N disk drives. Thus, if a single drive on a RAID device fails, the piece of data block that was stored on the failed drive can be restored/reassembled [col 1, lines 20-59]. Grun further discloses "descriptors" that are used to accomplish the actual data movement from host computer (10) to RAID device (40). Data transfer is initiated using VI RDMA transfer facility [col 3, lines 14-45] [Fig. 2].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of the adapter wherein the execution engines are adapted to generate multiple gather entries corresponding to a single one of the work items, *each of the gather entries defining no more than a single packet* to be generated by one of the gather engines, as disclosed by Grun, for the motivation of reducing data transmission delay and efficiently striping data for data redundancy [col 1, lines 33-46].

17. Claims 39 and 40 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tzeng in view of Snyder II et al, (hereinafter Snyder), U.S. Patent 6,888,830.

18. As per claim 39, Tzeng in view of Snyder II discloses an adapter according to claim 28, wherein each of the execution engines comprises: a buffer for holding the work items to be processed, the buffer having a programmable watermark level; an execute machine, coupled to read the work items from the buffer and to generate the gather entries corresponding thereto; and a fetch machine, coupled to fetch the work items from the memory to the buffer responsive to a volume of the work items in the buffer having fallen below the watermark level.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received

work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 28 comprising a memory interface, a first plurality of execution engines to generate gather entries (descriptors) defining packets to be transmitted over the network responsive to the work items, a scheduling processor, and switching circuitry to submit the gather entries to the gather engines for generation of the packets responsive thereto; he does not explicitly disclose the adapter wherein each of the execution engines comprises: a buffer for holding the work items to be processed, the buffer having a programmable watermark level; an execute machine, coupled to read the work items from the buffer and to generate the gather entries corresponding thereto; and a fetch machine, coupled to fetch the work items from the memory to the buffer

responsive to a volume of the work items in the buffer having fallen below the watermark level.

This feature is disclosed by Snyder in his invention of an integrated circuit that processes a communication packet and comprises a core processor and scheduling circuitry. The core processor executes a software application that directs the core processor to process the communication packet. The scheduling circuitry retrieves first scheduling parameters cached in a context buffer for the packet and executes a first algorithm based on the first scheduling parameters to schedule subsequent transmission of the communication packet [Abstract]. Snyder further discloses a buffer having programmable watermark level [col 8, 22-40], hardware engines performing sophisticated searches for channel identifiers and automatic fetching of context information from the buffers [col 5, line 58 – col 6, line 20] [col 7, line 55 – col 8, line 40] [Figs 1-3].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of the adapter wherein each of the execution engines comprises: a buffer for holding the work items to be processed, the buffer having a programmable watermark level; an execute machine, coupled to read the work items from the buffer and to generate the gather entries corresponding thereto; and a fetch machine, coupled to fetch the work items from the memory to the buffer responsive to a volume of the work items in the buffer having

fallen below the watermark level, as disclosed by Snyder, for the motivation of providing robust functionality in extremely high-speed packet streams, via a Stream Processor integrated circuit [col 1, lines 33-46].

19. As per claim 40. An adapter according to claim 39, wherein the watermark level is programmable responsive to the classes of service of the transport service instances assigned to each of the execution engines for generation of the corresponding gather entries.

Tzeng discloses as his invention a method of determining an order of received work queue entries based on respective service levels, and outputting the received work queue entries according to the determined order, the invention featuring a *host channel adapter comprising a pre-link module, a transport service module, and a post-link module*. The pre-link module is configured for determining an order of received work queue entries based on respective service levels, the pre-link module outputting the received work queue entries according to the determined order. The transport service module is configured for generating transport layer headers for the work queue entries output from the pre-link module according to the determined order, and the post-link module is configured for generating, in the determined order, transmit packets having the respective transport layer headers for output onto a network [col 2, line 46 – col 3, line 3].

While Tzeng discloses substantial features of the invention such as the adapter of claim 28 comprising a memory interface, a first plurality of execution engines to generate gather entries (descriptors) defining packets to be transmitted over the network responsive to the work items, a scheduling processor, and switching circuitry to submit the gather entries to the gather engines for generation of the packets responsive thereto; he does not explicitly disclose the adapter wherein each of the execution engines comprises: a buffer for holding the work items to be processed, the buffer having a programmable watermark level; an execute machine, coupled to read the work items from the buffer and to generate the gather entries corresponding thereto; and a fetch machine, coupled to fetch the work items from the memory to the buffer responsive to a volume of the work items in the buffer having fallen below the watermark level.

This feature is disclosed by Snyder in his invention of an integrated circuit that processes a communication packet and comprises a core processor and scheduling circuitry. The core processor executes a software application that directs the core processor to process the communication packet. The scheduling circuitry retrieves first scheduling parameters cached in a context buffer for the packet and executes a first algorithm based on the first scheduling parameters to schedule subsequent transmission of the communication packet [Abstract]. Snyder further discloses a buffer having programmable watermark level [col 8, 22-40], hardware engines performing sophisticated searches for channel identifiers and automatic fetching of context

information from the buffers [col 5, line 58 – col 6, line 20] [col 7, line 55 – col 8, line 40]
[Figs 1-3].

It would therefore be obvious to one of ordinary skill in the art at the time of the invention to modify and/or combine Tzeng's invention with the feature of the adapter wherein each of the execution engines comprises: a buffer for holding the work items to be processed, the buffer having a programmable watermark level; an execute machine, coupled to read the work items from the buffer and to generate the gather entries corresponding thereto; and a fetch machine, coupled to fetch the work items from the memory to the buffer responsive to a volume of the work items in the buffer having fallen below the watermark level, as disclosed by Snyder, for the motivation of providing robust functionality in extremely high-speed packet streams, via a Stream Processor integrated circuit [col 1, lines 33-46].

Conclusion

1. The Examiner has cited particular columns and line numbers in the references applied to the claims above for the convenience of the applicant. Although the specified citations are representative of the teachings of the art and are applied to specific limitations within the individual claim, other passages and figures may apply as well. It is respectfully requested from the applicant in preparing responses, to fully consider the

references in entirety as potentially teaching all or part of the claimed invention, as well as the context of the passage as taught by the prior art or disclosed by the Examiner.

2. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

- Haugseth et al, Patent No. 6,856,619
- Ben-Ze'Ev et al, Patent No. 6,687,757

3. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Glenford Madamba whose telephone number is 571-272-7989. The examiner can normally be reached on M-F 8:30-5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Zarni Maung can be reached on 571-272-3932. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should

you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Glenford Madamba
Examiner
Art Unit 2151



Patrice Winder
PATRICE WINDER
PRIMARY EXAMINER